

Claims

1. A method of controlling a queue buffer (20), said queue
buffer (20) being connected to a link (40) and being
5 arranged to queue data units (30) of a flow in a queue
(21), comprising

determining (S1) a value of a length parameter (QL;
QL_{av}) related to the length of said queue,

10 comparing (S2) said value (QL; QL_{av}) with a length
threshold value (L_{th}) and performing (S3) a congestion
notification procedure if said value (QL; QL_{av}) is
greater than said length threshold value (L_{th}), and

15 an automatic threshold adaptation procedure (S6),

characterized in that

20 said automatic threshold adaptation procedure (S6)
comprises a procedure for adjusting said length
threshold value (L_{th}) on the basis of one or more flow
control parameters (LIM1, LIM2; rwnd; rlfir).

- 25 2. The method of claim 1, wherein said one or more flow
control parameters (LIM1, LIM2) are predetermined
values.

3. The method of claim 2, wherein said predetermined values
30 are associated with known flow control procedures for
one or both of data unit senders and data unit
receivers.

4. The method of one of claims 1 to 3, furthermore
35 comprising a procedure (S5) for determining one or more
of said one or more flow control parameters (LIM1, LIM2;

rwnd; rlfr) from a flow control parameter (rwnd; rlfr) introduced by one of a sender and a receiver of said flow queued in said queue (21).

- 5 5. The method of claim 4, wherein said flow control parameter is introduced by said receiver and inserted into acknowledgment data units sent from said receiver to said sender for acknowledging the correct receipt of data units.
- 10 6. The method of claim 5, and said buffer being provided in a network node of a communication network connecting said sender and said receiver, wherein said procedure for determining said flow control parameter comprises
- 15 extracting said flow control parameter from said acknowledgement data units at said network node.
- 20 7. The method of claim 5 or 6, and said buffer being provided in a first network node of a communication network connecting said sender and said receiver, wherein said procedure for determining said flow control parameter comprises extracting said flow control parameter from said acknowledgement data units at a
- 25 second network node different from said first network node and sending said flow control parameter from said second network node to said first network node.
- 30 8. The method of one of claims 1 to 7, and a flow control performed for said flow in said queue (21) being window-based, wherein one of said one or more flow control parameters is a control window.
- 35 9. The method of claim 8, wherein said control window is introduced by said receiver and expresses a limitation of how many data units the receiver can handle.

10. The method of claim 8, wherein said control window is introduced by said sender and expresses a limitation of how many data units the sender can send.

5 11. The method of one of claims 1 to 7, and a flow control performed for said flow in said queue (21) being rate-based, wherein one of said one or more flow control parameters is a control rate.

10 12. The method of claim 10, wherein said control rate is introduced by said receiver and expresses a data rate limitation for arriving data units that the receiver can handle.

15 13. The method of claim 10, wherein said control rate is introduced by said sender and expresses one of a data rate limitation for the rate of data units that the sender can send, a current sending rate and a target sending rate.

20

14. The method of one of claims 1 to 13, wherein said automatic threshold adaptation procedure comprises

estimating a link capacity value (LC),

25

analyzing whether the performance of a congestion notification procedure will lead to an underutilization of said link due to a reaction of said sender to the congestion notification under the condition that said length threshold value (L_{th}) is set equal to said estimated link capacity value (LC), and

30

adapting said length threshold value (L_{th}) on the basis of a result of said analyzing step by setting said length threshold value (L_{th}) equal to said estimated link capacity value (LC) if said analyzing step indicates no underutilization, and setting said length

35

threshold value (L_{th}) larger than said estimated link capacity value (LC) otherwise.

15. The method of claim 14, wherein said length threshold value (L_{th}) is set to a value derived on the basis of one of said flow control parameters (LIM1, LIM2; rwnd; rlfir) if said analyzing step indicates underutilization.

16. The method of claim 14 or 15, and a sender of said flow in said queue (21) sending said data units in a predetermined sequence, a receiver of said flow in said queue (21) sending to said sender acknowledgment messages for acknowledging the correct receipt of said data units, where each acknowledgment message identifies the last data unit correctly received in said sequence, and said receiver sending to said sender a first window value expressing a limitation of how many data units the receiver can handle, a flow control performed by said sender being window-based using a send window, said send window being selected as the minimum of said first window value and a second window value, such that said sender must not send data units with a sequence number higher than the sum of the highest acknowledged sequence number and the send window, and said sender dividing said second window value by two as a reaction to a congestion notification, and thereafter increasing the second window by a predetermined increment for each duplicate acknowledgment message it receives,

wherein

one of said one or more flow control parameters (LIM1, LIM2; rwnd; rlfir) is said first window value and said length threshold value (L_{th}) is initially set equal to said estimated link capacity value (LC), and

said automatic threshold adaptation procedure comprises setting said length threshold value (L_{th}) equal to said estimated link capacity value (LC) if said first window value is greater than 1.5 times the sum of said
5 estimated link capacity value (LC) and the momentary value of said length threshold value (L_{th}).

17. The method of claim 16, wherein said automatic threshold adaptation procedure comprises setting said length
10 threshold value (L_{th}) equal to said estimated link capacity value (LC) if said first window value is greater or equal to 1.5 times the sum of said estimated link capacity value (LC) and the momentary value of said length threshold value (L_{th}).

18. The method of claim 16 or 17, wherein said length
15 threshold value (L_{th}) is set equal to a function of said first window value if said first window value does not fulfil the condition for setting said length threshold value (L_{th}) equal to said estimated link capacity value (LC).

19. The method of claim 18, wherein said function is the
20 difference between said first window value and a predetermined reduction value.

20. A computer program product comprising a computer program
25 arranged to execute the method of one of claims 1 to 19 when executed on a programmable data processing device connected to a communication network containing said link.

21. A queue buffer controller (10) for controlling a queue
30 buffer (20) that is connected to a link (40) and is arranged to queue data units (30) of a flow in a queue (21), comprising

a queue length determinator (101) for determining a value of a length parameter (QL ; QL_{av}) related to the length of said queue,

5 a comparator (102) for comparing said value with a length threshold value,

a congestion notifier (103) for performing a congestion notification procedure if said value is greater than
10 said length threshold value, and

a threshold adaptor (104) for automatically adapting said length threshold value,

15 characterized in that

said threshold adaptor (104) being arranged for adjusting said length threshold value on the basis of one or more flow control parameters ($LIM1$, $LIM2$; $rwnd$; $rlfr$).
20

22. The queue buffer controller (10) of claim 21, wherein said one or more flow control parameters ($LIM1$, $LIM2$) are predetermined values.

25

23. The queue buffer controller (10) of claim 22, wherein said predetermined values are stored in said queue buffer controller (10) and associated with known flow control procedures for data unit receivers.

30

24. The queue buffer controller (10) of one of claims 19 to 21, further comprising a flow control parameter determinator (105) for determining one or more of said one or more flow control parameters ($LIM1$, $LIM2$; $rwnd$; $rlfr$) from a flow control parameter ($rwnd$; $rlfr$) introduced by one of a sender and a receiver of said flow queued in said queue (21).
35

25. The queue buffer controller (10) of claim 24, wherein
said flow control parameter is introduced by said
receiver and inserted into acknowledgment data units
sent from said receiver to said sender for acknowledging
the correct receipt of data units, and said queue buffer
(20) being provided in a network node of a communication
network (3) connecting said sender and said receiver,
wherein said flow control parameter determinator (105)
is arranged for extracting said flow control parameter
from said acknowledgement data units at said network
node.

26. The queue buffer controller (10) of claim 21, wherein
said flow control parameter is introduced by said
receiver and inserted into acknowledgment data units
sent from said receiver to said sender for acknowledging
the correct receipt of data units, and wherein said
queue buffer (20) is provided in a first network node of
a communication network (3) connecting said sender and
said receiver, wherein said flow control parameter
determinator (105) is arranged for receiving said flow
control parameter from a second network node at which
said flow control parameter was extracted.

27. The queue buffer controller (10) of one of claims 21 to
26, wherein said threshold adaptor (104) comprises

an estimator for estimating a link capacity value (LC),

an analyzer for analyzing whether the performance of a
congestion notification procedure will lead to an
underutilization of said link due to a reaction of said
sender to the congestion notification under the
condition that said length threshold value is set equal
to said estimated link capacity value (LC), and

an adaptor for adapting said length threshold value on the basis of a result of said analyzing step by setting said length threshold value equal to said estimated link capacity value (LC) if said analyzing step indicates no underutilization, and setting said length threshold value larger than said estimated link capacity value (LC) otherwise.

28. The queue buffer controller (10) of claim 27, wherein said threshold adaptor (104) is arranged to set said length threshold value to a value derived on the basis of one of said flow control parameters (LIM1, LIM2; rwnd; rlfr) if said analyzer indicates underutilization.

29. The queue buffer controller (10) of claim 27 or 28, a sender of said flow in said queue sending said data units in a predetermined sequence, a receiver of said flow in said queue sending to said sender acknowledgment messages for acknowledging the correct receipt of said data units, where each acknowledgment message identifies the last data unit correctly received in said sequence, and said receiver sending to said sender a first window value expressing a limitation of how many data units the receiver can handle, a flow control performed by said sender being window-based using a send window, said send window being selected as the minimum of said first window value and a second window value, such that said sender must not send data units with a sequence number higher than the sum of the highest acknowledged sequence number and the send window, and said sender dividing said second window value by two as a reaction to a congestion notification, and thereafter increasing the second window by a predetermined increment for each duplicate acknowledgment message it receives,

wherein

one of said one or more flow control parameters (LIM1, LIM2; rwnd; rlfir) is said first window value and said threshold adaptor (104) is arranged to initially set said length threshold value equal to said estimated link capacity value, and to set said length threshold value equal to said estimated link capacity value if said first window value is greater than 1.5 times the sum of said estimated link capacity value (LC) and the momentary value of said length threshold value (Lth).

30. The queue buffer controller (10) of claim 29, wherein said threshold adaptor (104) is arranged for setting said length threshold value equal to said estimated link capacity value if said first window value is greater or equal to 1.5 times the sum of said estimated link capacity value (LC) and the momentary value of said length threshold value (Lth).

31. The queue buffer controller (10) of claim 29 or 30, wherein said threshold adaptor (104) is arranged to set said length threshold value equal to a function of said first window value if said first window value does not fulfil the condition for setting said length threshold value equal to said estimated link capacity value.

32. The queue buffer controller (10) of claim 31, wherein said function is the difference between said first window value and a predetermined reduction value.